

METHOD AND SYSTEM FOR PARAMETERIZED WEB DOCUMENTS

Inventors:

Balas Natarajan Kausik, Ph.D.
Janardhanan Jawahar

Prepared By:

Joseph Yang, Ph.D.
Skadden, Arps, Slate, Meagher & Flom LLP
525 University Avenue
Palo Alto, California 94301
(650) 470-4500

I hereby certify that this correspondence is being deposited with the United States Postal Service as "Express Mail Post Office to Addressee" under 37 CFR § 1.10 (Label No. EL 728 499 948 US) in an envelope addressed to: Commissioner for Patents, Box Patent Application, Washington, D.C. 20231, on April 19, 2001.

Jan Steele

METHOD AND SYSTEM FOR PARAMETERIZED WEB DOCUMENTS

CROSS-REFERENCE TO RELATED DOCUMENTS

5

This application is a continuation-in-part of co-pending application Serial No. 09/634,134, filed on August 8, 2000, which is hereby incorporated by reference in its entirety. This application is related to co-pending application Serial No. 09/816,802, filed on March 23, 2001.

10

FIELD OF THE INVENTION

The present invention relates to accelerating the delivery of content and reducing congestion in a networked environment. In particular, a second document can be described as a modification of one or more first documents in such a manner that the second document can be downloaded and correctly displayed by commonly deployed content browsers without necessarily requiring additional software.

BACKGROUND

20

It is commonly required to personalize a web document for each user who views the document. Such personalization might involve an advertisement that is targeted on the basis of the preferences of the user, or information such as quotes for stocks in the user's portfolio etc. Furthermore, it is also commonly required to update pages that a user or a group of users has previously viewed with fresh information on stock prices, news etc. In either case, the web document requested by the user comprises content that is common to another web document previously delivered to the same user or other users, as well as content that is new and possibly particularized to the user. Furthermore, the common content typically forms the majority of the bytes in the document. In other words, only a small percentage of the document changes between subsequent downloads.

In the background art (see, e.g., U.S. Patent 6,178,461), it is known to encode web documents as variations of previously delivered documents that require explicit action on

the part of the user's content browser. In particular, the user's content browser selects one or more objects from its cache that it expects to be similar to the current document and includes references to these documents when requesting the current document from the content server. The difficulty with this browser-driven approach is that it requires millions of content browsers to be upgraded to include this modification. Furthermore, it requires that the server have access to the millions of old documents previously transmitted to the users in order to correctly recover the base documents referenced in each request.

Another approach was disclosed in co-pending U.S. patent application Serial No. 09/634,134, filed on August 8, 2000, which is hereby incorporated by reference in its entirety. That application disclosed general techniques for transmitting the incremental differences between successive downloads of web documents were disclosed. In the present application, we expand upon those teachings.

15

SUMMARY

The present invention includes methods and systems for constructing web (or other networked) documents as parameterized forms of other web (or other networked) documents. For example, a document may be represented as a collection of changes and insertions to be applied to one or more first documents, where the first document(s) is incorporated by reference in the second document. Such first documents are typically previously delivered documents that may be in the local cache of the user's content browser, or a network cache common to several users. Thus, instead of delivering the entire document over the slower wide-area network connecting the content server to the content browser, the document is delivered as a collection of changes to previously delivered documents that are much closer to the user's content browser. For example, the collection of changes may travel across a wide-area network, while the first documents are accessed from the local cache of the browser or from a network cache across a local-area network. The foregoing has the advantages of reducing both bandwidth usage as well as the time required to deliver a document, and is particularly well suited for (but not

limited to) dynamically generated and personalized content. Examples of dynamically generated, personalized, and continually changing content are stock quotes, account information, personalized news feeds, etc.

In an exemplary embodiment of the invention, special software is not necessarily required at the end user for reconstructing the second document from the set of base documents and the set of modifications to the base documents. This is of commercial significance where distributing special software to millions of users may be an obstacle. In various aspects of this exemplary embodiment, the server may decide which documents to use as base documents, and may also maintain copies of those documents it intends to use as base documents, in a manageable and controlled fashion. Lastly, since these base documents will be frequently referenced in many requests, they may also be stored in network caches.

DETAILED DESCRIPTION

For convenience, the invention will be described herein with respect to "documents" (or, equivalently, "files") which should be understood to include any content-bearing items transmitted in a networked environment, including, without limitation, text, data or graphic files (or combinations thereof), images, objects, programs, scripts, audio, video, and other kinds of documents. More specifically, one preferred embodiment of the invention is described with respect to documents comprised of byte strings. Those skilled in the art will understand that the teachings of the invention readily extend to other forms of documents deliverable on a network. Thus, the term "strings" should be understood to be equally applicable to other types of content elements as appropriate to the nature of the document.

The system contemplates a computer and software running thereon. The system takes as input a "current" document, which is to be transmitted to the user. The system then selects one or more first documents, which are the "base" documents to be incorporated by reference in the transmission. Of course, a "current document" is not necessarily the latest available version as of the time of transmission, only that it supersedes a base version of the document. Similarly, a "base document" is not

necessarily the earliest available document, or even one that has actually been sent to a particular user (e.g., a base document could be a template stored on the content server), only that it forms the basis for the "current document" to be transmitted to the user.

The base document is typically selected on the basis of its similarity to the current
5 document, and is typically selected to be an older and previous version of the same document or related document. For example, if the current document is a brokerage report on a particular stock, the base document could be an older report on the same stock, or an older report on a different stock. As another example involving online retail, if the current document describes an item of clothing, the base document could be
10 describing another related item. By comparing the base documents with the current document, the current document is decomposed into strings that occur in one of the base documents and strings that do not occur in any of the base documents. Techniques for efficient comparison of the base and current documents to identify the various substrings are disclosed in co-pending U.S. patent application 09/634,134. The current document is
15 then represented as a series of substrings of the base documents, interspersed with clear-text strings that do not occur in any of the base documents. The representation is encoded as a program in a scripting language such as Javascript that can be readily executed by common content browsers, and upon execution, causes the current document to be reconstructed and displayed by the content browser.

20 The base document is typically selected with respect to the context of the user's request for the current document. One possibility is to set the base document to be an older version of the current document, and to periodically update the base document when the size of the changes between the current document and the base document exceeds a certain limit. Another possibility is to dynamically select the base document as
25 a central and representative document from amongst a collection of documents. In yet another possibility, the base document could be a template document explicitly constructed for the purpose, and never delivered in a visible form to the content browser. More detail about these possibilities are disclosed in co-pending U.S. patent application Serial No. 09/816,802, filed on March 23, 2001, which is hereby incorporated by
30 reference in its entirety.

Example of Operation Involving Text Documents

As an example of operation, consider the situation where the documents are text
5 documents. Further suppose that the base document consists of the text “pack my box
with five dozen liquor jugs”, and the second document consists of the text “pack my box
with five dozen liquor mugs.” It is clear that characters 1 through 35 and characters 37
through 40 of the second document are same as those of the base document. The two
documents differ only at character position 36, where the character "m" occurs in the
10 second document as compared to the character "j" in the base document. The proposed
system constructs the following programmatic representation of the second document that
is exemplary of programs in scripting languages such as Javascript, which are
“interpretable” by the content browser in that programs written in such languages can be
executed on-the-fly by the content browser. Those skilled in the art will realize that this
15 can also be implemented in languages that are executed via being compiled rather than
via being interpreted.

```
var base_string = base_document;  
print(base_string,1,35);  
20    print("m");  
print(base_string,37,40);
```

When downloaded and executed by a content browser, the above program will
display characters 1 through 35 of the base document, the clear text character “m” and
25 then characters 37 through 40 of the base document.

Other Embodiments and Aspects of the Invention

If the base document has been previously downloaded by the content browser and
30 is resident on the browser’s cache, the browser may use the cached copy, saving the time

and bandwidth required to download the base document. If the base document is not in the browser cache, the browser may request the base document from the network. For example, in an exemplary embodiment of the invention, as the browser's request for the base document travels towards the content server, it may encounter one or more network
5 caches or "proxies" at points of aggregation such as the enterprise network point or the Internet Service Provider's network head. A network cache or other point of aggregation containing the base document can respond to the browser's request therefor. Thus, the content server can present documents as dynamic updates to previously transmitted and cached versions of other documents, without explicit regard to where such base
10 documents may be cached.

In order to enable that base documents be available for repeated use over long periods of time, it is beneficial to configure them to bear names or URLs that are unlikely to be repeated or conflict with the URLs of other documents. For example, the base documents might be assigned URLs that are randomly selected integer IDs. By making
15 these IDs sufficiently long, the likelihood that such a URL will clash with another can be made impractically small. Still other ways for minimizing conflicts between names will be apparent to those of skill in the art, and need not be described in detail here.

Another consideration is the life of a base document in a network or browser cache. Such caches typically require that the valid life of a document be declared
20 explicitly at the time of transmission. For example, a network cache will continue to use a base document over the specified life of the document. Once that life has expired, the network cache will discard that document. In order to reuse base documents over the longest life, it is beneficial that base documents carry a lifetime that is greater than their expected usage time. Their expected life can typically be estimated by the average time
25 over which the difference between the current document and the base document exceeds some preset limit (at which point the base document is typically replaced).

The various embodiments described above should be considered as merely illustrative of various embodiments of the present invention. Those skilled in the art will realize that the present invention in its most general form is applicable regardless of
30 whether the user is a person at a networked computer or a wireless device, or a network

device such as a cache or proxy agent. Those skilled in the art will also realize that various aspects of the present invention are applicable to the full range of data forms transmitted on the Internet or other types of networks, including but not limited to text, images, video and audio.

5 For example, consider a situation wherein an image encoded in JPEG is to be delivered to the user's content browser, with a small modification that is particular to the user. In this situation, the base image is the original image, and the current image is the modified image to be delivered. Since JPEG encoded images are treated typically as 8x8 blocks of pixels, the modified image can be described as a parameterization of the base
10 image, wherein the modified blocks are to replace the original blocks. That is, blocks in JPEG images are equivalent to "strings" of a document.

Likewise, if the image were to be encoded in GIF format which uses run length encoding, the image is described as a series of runs of the same color. A modification of such an image can be described as a parameterization, wherein certain runs of the base
15 image are to be replaced by certain other runs of the modified image. That is, runs in a GIF image are equivalent to "strings" of a document.

As another example, consider a situation wherein a sequence of digitized video is to be delivered to the user's content browser (e.g., a video player or decoder), but that the sequence of video is to be interspersed with some advertising that is targeted to each specific user. In this situation, the "base sequence" (or base document) is the original piece of video, and the "current sequence" (or current document) is the piece of video with the advertisements inserted. Those skilled in the art will immediately realize that the current sequence can be described as a parameterization of the base sequence, wherein the advertisements are specified to be inserted at specific timing points in the
20 base sequence. More specifically, consider a video sequence embedded in the MPEG format. Typically, the MPEG format bundles approximately 15 frames of video into a block, and treats each block independently of the other blocks. At normal playback, 15 frames represent one half of one second in elapsed time. Since each block of frames is independent of other blocks, video sequences representing advertisements and other
25 customized material can be interspersed between two blocks as a third independent block.
30

That is, blocks of frames in MPEG images are equivalent to "strings" of a document. The example below shows one embodiment of the invention for video sequences, on a hypothetical MPEG player. Specifically, the example depicts the insertion of an advertisement sequence between frames 450 and 451 of the original sequence.

5

```
var base_MPEG_sequence = movie;  
play (base_MPEG_sequence,1,450);  
play (advertisement_sequence_1);  
play (base_MPEG_sequence,451,1500);
```

10

Those skilled in the art will realize that digitized audio sequences in MP3 and similar formats also use block-based encoding as in MPEG, can be treated in a fashion similar to digital video.

15

Thus the various embodiments and aspects described above are not intended to be exhaustive or to limit the invention to the forms disclosed. Those skilled in the art will readily appreciate that still other variations and modifications may be practiced without departing from the general spirit of the invention set forth herein. Therefore, it is intended that the invention be defined by the claims that follow.